

# EMBEDDED x86:

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR



X86 MICROPROCESSORS MAY NOT HAVE THE ABSOLUTE LOWEST POWER PER MEGAHERTZ, ABSOLUTE HIGHEST PERFORMANCE PER CLOCK TICK, OR OPTIMUM CODE EFFICIENCY WHEN YOU EVALUATE THEM AGAINST ALL POSSIBLE ALTERNATIVES. BUNDLE UP THE FEATURES, THOUGH, AND THE TOTAL "KIT" OFTEN STACKS UP FAVORABLY VERSUS THE COMPETITION.

# KEYSTONE OF YOUR NON-PC DESIGN?

**T**hose of you considering PC building blocks for your non-PC designs would do well to keep in mind, as your counterparts who are veterans of this architecture direction have already learned, that placing your stakes on a PC roll of the dice isn't a sure bet. On the one hand, you'll benefit from the tremendous pace of innovation endemic to the PC business, along with low prices resulting from the hundreds of millions of PCs sold worldwide each year. On the other hand, that same tremendous pace of innovation also translates to a tremendous pace of obsolescence, which can be problematic for systems with production cycles that measure longer than six months!

Assuming that you build enough sourcing flexibility into your design to comprehend supply impermanence, PC-tailored microprocessors can provide a cost-effective means of achieving your system's performance targets. And, with the PC industry's amplified power-consumption focus, battery life, power-supply size, and heat dissipation aren't the concerns they might have previously been, either. Traditional CPU and DSP suppliers haven't stood still in the face of the PC-processor onslaught, however, and their alternative solutions remain optimal in many situations. A solid understanding of the historical trends, current status, and future plans of the primary x86 CPU suppliers will enable you to assess which path to take for your next design (see **sidebar** "Montalv-who?").

## **INTEL: BACK ON TRACK**

Intel exemplifies how substantially a company's fortunes can change in five short years. In the early portion of this decade, Intel based its entire microprocessor product line (laptops to servers), with the exception of the Itanium processor, on the NetBurst microarchitecture (see **sidebar** "Speeds and feeds"). NetBurst had lengthy pipelines—20 stages in the initial 180-nm Willamette variant, extending to as many as 31

stages in the final 90-nm Prescott and 65-nm Cedar Mill iterations. These pipelines performed well when code characteristics paired them with highly predictable multimedia-instruction streams. But the low IPC (instructions-per-clock) attribute inherent in any long-pipeline approach, combined with substantial branch-misprediction penalties, gave NetBurst underwhelming performance on more conventional code. And, in striving to boost clock rates as compensation for long-pipeline penalties, Intel ran into substantial leakage-current problems beginning at the 90-nm process node, which rendered the company's NetBurst products 6.2 GHz short of the initial 10-GHz microarchitecture target. (Even getting to 3.8 GHz proved to be a formidable project.)

Intel's fortunes began trending back upward, beginning in the mobile-computing segment, when, in the spring of 2003, it introduced the first Banias iteration of the Pentium M microarchitecture (**Reference 1**). Fabricated on a 130-nm process, Banias preceded the 90-nm-based Dothan with a larger L2 cache and 65-nm dual-core Yonah with a shared L2 cache (**Reference 2**). Pentium M leveraged and expanded on the execution unit of the Pentium III and coupled it to the Pentium 4 bus interface. As such, it offered more efficient power and instructions than NetBurst on a per-clock basis with conventional code traces. Yonah-generation CPUs, instead of using the Pentium M brand of their predecessors, employed a Core marketing moniker that proved somewhat confusing a short time later, when the company rolled out a suite of 65-nm-process-based Merom, Conroe, and Wolfdale CPUs, respectively spanning laptops to servers, and leveraging the follow-on Core microarchitecture, which the company marketed with the Core 2 promotion brand.


Intel is now shipping the second iteration of its Core microarchitecture, known as Penryn, which it fabricates on a 45-nm-process lithography. Penryn reflects the company's "tick-tock" strategy, a product cadence that involves smaller lithography products, with only minor corresponding feature tweaks (tick), followed roughly one year later by a more substantive architecture revamp on a common process foundation (tock).

As such, the corresponding tock to today's Penryn tick, Nehalem, ramps into production this year, and Intel is publicly demonstrating it in prototype-system form. Nehalem will address several longstanding AMD criticisms, albeit ones that few benchmarking tests to date have shown result in real-life performance shortcomings. Through today's Penryn products, all intercore communication, with the exception of intradie shared-cache-coherency synchronization—whether within a


die, between dice in a multichip monolithic-packaged CPU, or between packaged CPUs—occurs through the same front-side bus that carries data traffic to and from external subsystems. The primary external subsystem is the core-logic chip set, which in today's designs contains the DRAM controller.

## MANY CAPABILITIES

Nehalem-class CPUs integrate the dedicated QuickPath Interconnect interprocessor—formerly, CSI (common-system interface). The link is conceptually reminiscent of the HyperTransport link, which AMD introduced in 2001 for communications between multiple cores on a die and between multiple-die and packaged CPUs. Also reminiscent of technology AMD pioneered in 2003 with the Athlon 64 and Opteron K8 (also known as Hammer) CPUs, Nehalem-based products embed DRAM controllers to, among other things, reduce the extended latencies systems now experience when cache misses require external-memory accesses. And speaking of cache, whereas today's Intel products go beyond two cores by combining multiple die under a common package lid, the prodigious transistor budget that the 45-nm process affords will enable the company to monolithically squeeze at least six CPU cores onto a single sliver of Nehalem-based Dunnington silicon (**Figure 1**). Each pair of cores, like other current products, shares a common L2 cache, and all six cores split a common L3 cache in the layout hole in which a fourth two-



(a)



(b)

**Figure 1** An advanced 45-nm-process technology enables Intel to both build ultra-high-transistor-count ICs, such as the upcoming six-core-plus-L3-cache Dunnington CPU (a), and ultrasmall chips, such as the first-generation Atom microprocessor (b).

## AT A GLANCE

▶ Rapid feature improvements and low prices make x86 CPUs appealing in many non-PC designs.

▶ After hitting a market-share low point in 2005, Intel has vigorously recovered over the past few years.

▶ Intel's setbacks boosted AMD's early-decade fortunes; the two companies have subsequently swapped success and failure stories.

▶ Via Technologies pioneered cost and power concentrations that other x86 suppliers later also embraced. A superscalar out-of-order CPU is finally on the way.

core cluster might otherwise go.

The 45-nm-process generation enables Intel to build not only cost-effective large-die products, but also very cost-effective small-die processors. That cost-effectiveness is the impetus for the Atom-CPU-product line, which Intel formally introduced at the Shanghai Intel Developer Forum in early April (**Reference 3**). Formerly Silverthorne, Atom combines with a single-chip companion device; Intel previously referred to the chip set as Menlow. Atom's origins derive from the under-development and x86-based Larrabee PC coprocessor, intended for graphics, imaging, physics, and other functions. Intel's architects determined that they required the ability to cost-effectively embed 16 or more x86 cores on a single Larrabee die and that the out-of-order execution and other exotic attributes of

the company's mainstream CPUs represented overkill for the targeted applications. Consequently, Intel went "back to the future," dusting off its Pentium III schematics to come up with an area-optimized CPU-core design for Larrabee. The company is attempting to maximize its return on investment by also developing few-physical-core chips that it bases on the Larrabee atomic building block, some also with HyperThreading virtual-multicore support, for power- and cost-sensitive mobile systems.


First-generation Atom CPUs come in five versions, with clock speeds that reach 1.86 GHz and a TDP (thermal-design-power) range of 0.65 to 2.4W. Corresponding average- and idle-power ranges are, respectively, 160 to 220 mW and 80 to 100 mW. The partner system-controller hub, available in three versions, features a 3-D graphics core; a hardware-accelerated, high-definition-video-decoding engine; high-definition-audio processing; and support for PCI (peripheral-component-interconnect) Express, USB, and SDIO (secure-digital-input/output) connectivity. And, with long-life-cycle embedded-system designs in mind, Intel promises at least seven years of product support. A planned dual-core Atom variant will be more compelling in low-cost laptop and desktop systems, and Intel also plans an even more integrated single-chip, albeit perhaps multichip, Moorestown Atom family for next year. All in all, after years of stumbling, Intel's seemingly back in full stride. Perhaps the biggest question on the company's road map for the remainder of the decade is the degree to which Atom will cannibalize Intel products in a manner that is fiscally unattractive to Intel, instead of broadening the overall x86 market at the expense of competitors, such as ARM, as Intel hopes.

## AMD: DESTINATION UNKNOWN

While key competitor Intel struggled through more than the first half of this decade, AMD "made hay," as the saying goes. The company's K7 Athlon microprocessor, which it introduced in 1999, proved to be a more conventional architecture than the NetBurst-based Pentium 4 Intel unveiled one year later and, as such, was highly clock-efficient from both performance and power-consump-

tion standpoints. Whereas Intel for many years attempted—largely without success, except in ultra-high-end configurations—to propel the 64-bit-system market to its proprietary, revolutionary Itanium processor, AMD chose a more evolutionary path that appended 64-bit instruction support onto the Athlon foundation. The result was 2003's K8-based Athlon 64 and Opteron (Reference 4). K8 CPUs provided other key evolutionary enhancements, as well, such as the earlier-mentioned HyperTransport links and integrated system-memory controllers. And AMD was also first to market with multicore x86 CPUs, at least from a monolithic-die standpoint, judging from 2005's multicore Opteron and Athlon 64 X2 introductions.

The last few years have, however, revealed AMD's key weakness: The company is a much smaller x86 player than Intel, both in employee-head-count and market-share metrics, so the success of each project it tackles is comparatively critical. AMD in 2003 began discussing the architectural goals for the K10 microarchitecture, and the company in



**Figure 2** AMD's CPUs' integrated memory controllers and intercore HyperTransport links represent unique capabilities, at least until Intel's Nehalem ramps into production this year.

2006 unveiled key details of its 65-nm-process-targeted monolithic quad-core Opteron, Barcelona, and Athlon-follow-on, Phenom. AMD ended up delaying Barcelona's introduction until September 2007 and then further postponed full-production shipments until two months ago, after the company fixed an embarrassing L2 cache TLB (translation-look-aside-buffer) flaw it discovered after Barcelona's launch and

coincident with the Phenom unveiling.

AMD shipped some Phenom material before the TLB fix, accompanying it with a BIOS-based microcode patch that had the unfortunate side effect of tangibly decreasing performance on many benchmarks. And, in part a reflection of AMD's being one process generation behind Intel, leading-edge Opteron and Phenom CPUs run at core-clock-speed and benchmark deficits compared with their Intel counterparts. On the one hand, this scenario is familiar to AMD; as its mid-1990s P (performance)-rating system suggests, the company's products have long exhibited lower clock speeds but more efficient clock usage than Intel offerings. However, whereas in the early part of this decade AMD was well-matched against Intel's NetBurst-based CPUs, it's now competing against the superior Core-derived follow-ons. AMD's processors' HyperTransport links and integrated DRAM controllers make up some of the clock-speed shortfall, but they can't completely bridge the gap. And, by the time AMD forecasts bringing 45-nm-based K10 follow-ons to market, Intel predicts that it will be ramping its first 32-nm Westmere CPUs into production, thereby preserving its one-generation lithography advantage.

The news, although troubling, isn't all bad on the AMD front. The company's aggressive pricing, at unknown per-product-profit impact, has enabled it to remain largely competitive with Intel in market areas in which its silicon's performance and power-consumption characteristics enable it to have a tangible presence. By virtue of its HyperTransport-based core-interconnect approach, along with other hooks that the K10 design included from the start, AMD can ship some partially defective dice as triple-core Phenom chips, thereby maximizing revenue per wafer (Figure 2). AMD is also moving its K8-based dual-core CPUs to the 65-nm process, which should improve their die-per-wafer usage and therefore increase AMD's combination of dedicated and foundry-supplied fabrication capacity.

Like Intel, AMD sells some of its products in embedded flavors with extended-life-cycle guarantees, enhanced testing, and other attributes that non-PC customers value. And AMD also continues


## MONTALVO-WHO?

In 1995, start-up Transmeta began working on a then-secret chip design. Over time, word leaked out that the company was readying an Intel-targeting, low-power x86 processor for portable-system applications, and, by the time Transmeta formally unveiled its plans in January 2000, public anticipation had reached fever pitch. Transmeta produced two product lines, Crusoe and Efficeon. Although Transmeta is now a shadow of its former self, having jettisoned the bulk of its staff and transformed itself into an IP (intellectual-property)-licensing entity, its industry influence lives on in the power-stingy chips that Intel and its competitors now sell.

Similarly secretive, several-year-old start-up Montalvo Systems may be following in Transmeta's footsteps. Basing opinions solely on tantalizing rumor tidbits and a few filed patents, industry observers believe that the company is readying a low-

power-x86-compatible-processor line based on an asymmetrical, function-specific, multicore architecture. To date, the company has raised more than \$73 million in private-equity funding, and its staff includes such long-term veterans of the Intel wars as Matt Perry, Transmeta's former and Montalvo's current chief executive officer, and Vinod Dham, the former vice president of Nexgen, former Intel chief architect, and former AMD vice president.

Information at press time suggests that Montalvo is running low on cash and seeking either additional private investment or a buyout. Will Montalvo ever turn its microprocessor vision into silicon, and, if so, will it be successful? There's more than one way to profitably return a backer's financial investment; Transmeta, for example, successfully concluded in late 2007 a 2006 lawsuit against Intel to the tune of \$250 million.



**Figure 3** Pico-ITX represents the newest and smallest iteration of Via's embedded-system-board series (a). The company scored a major volume-design coup when HP in early April introduced the C7-based model 2133 Mini-Note PC (b). But will HP base its next iteration of the Mini-Note platform on the C7, Via's next-generation Isaiah CPU (c) or Intel's Atom?

to sell the Geode integrated-x86 line, which it acquired from National Semiconductor in 2003 and whose legacy extends back to the Cyrix MediaGX. (National Semiconductor purchased Cyrix in 1997.) Although AMD doesn't officially comment on future plans for Geode, the most recently unveiled product variant, Geode LX, dates from 2005, and the company shut down its Geode-design center in 2006.

### VIA: SCRAPPY UNDERDOG

Via Technologies, long known as a core-logic-chip-set supplier, also became a CPU manufacturer in September 1999, when it acquired IDT's Centaur Technology subsidiary. (Via also purchased National Semiconductor's PC-specific Cyrix assets in September 1999.) From its beginnings, Centaur has consistently focused on minimizing die size and power consumption and striving to deliver sufficient performance for conventional applications. This stance was controversial in the company's early years, when a "clock-rate-is-everything" mentality dominated the industry. However, such thinking has become more mainstream.

The company's series of C3-generation CPUs, known in fanless configurations as Esther, today predominantly derives from the Nehemiah core, though Via bases some C3 products on the older

Samuel 2 design (Reference 5). Nehemiah features various general-architecture enhancements, such as a full-speed floating-point unit, a deeper pipeline, a more comprehensive MMX (multimedia-extensions) implementation, and a switch from 3DNow! to SSE (streaming-single-instruction/multiple-data-extensions)-enhanced multimedia-instruction sets. Nehemiah also targets embedded-system designs, such as twin thermal-based hardware RNGs (random-number generators). Via refers to these RNGs as PadLock RNGs and, in a later proliferation, an AES (Advanced Encryption Standard) engine, PadLock ACE. Single-chip CoreFusion products combining Nehemiah CPU and north-

bridge core-logic die in a unified package are also available.

For more conventional desktop and laptop designs, such as Hewlett-Packard's recently introduced model 2133 Mini-Note PC, along with higher end embedded-system applications with and without fans, Via also offers its C7-class CPUs (Figure 3). The company built them on partner IBM's 90-nm SOI (silicon-on-insulator) process, and they offer enhanced PadLock capabilities. As Via's documentation describes, "The Via PadLock Security Engine in the Via C7 processor adds SHA [secure-hash-algorithm]-1 and SHA-256 hashing for secure-message digests, and a hardware-based montgomery multiplier supporting key sizes up to 32 [kbytes] in length to accelerate public-key cryptography, such as RSA [Rivest/Shamir/Adleman]. The Via C7 also provides NX execute protection, providing protection from malicious software, such as worms and viruses, and is used in Microsoft Windows XP with SP2. Integrating security directly onto the processor die ensures speeds and efficiency many times that available in software, yet with negligible impact on processor performance."

Reflecting the company's core-logic heritage, a range of companion chip sets for Via's CPU families provides system-design flexibility, with integrated options such as 2- and 3-D graphics acceleration and hardware-accelerated decoding of MPEG-2, MPEG-4, and other video codecs. External-interface options include single-data-rate SDRAM, double-data-rate SDRAM in multiple generations, parallel- and serial-ATA (advanced-technology-attachment) mass storage, PCI and PCI Express add-in

## SPEEDS AND FEEDS

This article contains a large number of processor-project and -product names without providing detailed information on most of them. The omission was intentional for two key reasons. First, available page count is limited, and I've chosen to be comprehensive of vendors, architectures, and product lines rather than the specifics of any product. Second, many of the products in this article are not yet in production, and specifications will inevitably evolve, rendering obsolete some of the data I might include. The good news is that EDN's Web site provides a timely and unlimited supplement to this article. Hit Google, Zibb, or your favorite search engine to peruse EDN.com for already-published write-ups and keep an eye on the *Brian's Brain* blog for article addendums, identifiable by their "Embedded x86" tag lines.



⊕ See the "Embedded x86" posts at [www.edn.com/briansbrain](http://www.edn.com/briansbrain) for supplemental information on this article's topics.

⊕ Go to [www.edn.com/080529cs](http://www.edn.com/080529cs) for more information on the vendors this article mentions. You can also click on Feedback Loop to post a comment on this article.

cards, and various USB-port counts.

The 21×21-mm NanoBGA2 package common to many C3 and C7 CPU proliferations efficiently uses available system-board real estate. Speaking of system boards, if you'd prefer not to design your own, Via will happily sell you one of its 6.7×6.7-in. (170×170-mm) mini-ITX, 4.7×4.7-in. (120×120-mm) nano-ITX, and 3.9×2.8-in. (100×72-mm) pico-ITX boards. Effectively shrunk PC motherboards, they come in a diversity of CPU, core-logic chip-set, and support-chip combinations, including 10/100-Mbit Ethernet and GbE (gigabit-Ethernet) transceivers, greater-than-two-channel surround-sound analog and digital audio, TPMs (trusted-platform modules), IEEE-1394-interface ICs, and analog-video encoders, along with DVI (digital-video-interface) and LVDS (low-voltage-differential-signaling) digital-video outputs.

The mini-ITX form factor, which Via unveiled in 2001, is the most mature of the three options, and the company's success in promoting it as an industry standard has attracted competitive attention. Third-party partners sell mini-ITX boards based on both AMD and Intel CPUs, and Intel also manufactures its own mini-ITX products. Notably, the company has shown several Atom-based mini-ITX designs at numerous public forums in recent months.

And speaking of Atom, Via's competitive CPU response, Isaiah, is key to the company's future viability. Ironically, as Intel has stripped out-of-order execution and other superfluous features from its previous architectures, thereby resulting in an approach somewhat reminiscent of Centaur's nearly decade-old vision, Glenn Henry, president of Via's Centaur subsidiary, and his design team are poised to unleash Via's first three-way superscalar out-of-order architecture. The company formally unveiled Isaiah in late January. It also adds support for 64-bit instructions, hardware virtualization, and other much-needed features of modern competitive CPUs. Via will initially fabricate Isaiah on a 65-nm foundry lithography, in which design targets include doubling the integer performance and quadrupling the floating-point performance over that of the Via C7 at equivalent clock speed.

Isaiah's 65-nm power-consumption target is a 25W TDP at 2 GHz. Via is currently scheduling first Isaiah-based products for production in the second quarter, marking a one-quarter slip from the original publicly stated schedule. The company should be shipping the chips by the time you read this or soon thereafter. Via will provide companion core-logic chip sets for the CPU. And, in an interesting partnership perhaps reflecting the company's shrinking market share in the AMD- and Intel-targeted core-logic markets, Nvidia will sell Isaiah-tailored chip sets, too.**EDN**

## REFERENCES

- 1 Dipert, Brian, "And then there was one: one more time," *EDN*, April 29, 2004, pg 29, [www.edn.com/article/CA411162](http://www.edn.com/article/CA411162).
- 2 Dipert, Brian, "Double take: Reassessing x86 CPUs in embedded-system applications," *EDN*, April 27, 2006, pg 69, [www.edn.com/article/CA6325588](http://www.edn.com/article/CA6325588).
- 3 Dipert, Brian, "The Intel Developer Forum: ARM-wrestling," *EDN*, Sept 21, 2007, [www.edn.com/article/CA6480629](http://www.edn.com/article/CA6480629).
- 4 Dipert, Brian, "Itanium: 'Itanic' or full steam ahead?" *EDN*, April 27, 2006, pg 106, [www.edn.com/article/CA6325611](http://www.edn.com/article/CA6325611).
- 5 Dipert, Brian, "Hands-on project: And then there was one," *EDN*, March 4, 2004, pg 44, [www.edn.com/article/CA382777](http://www.edn.com/article/CA382777).

You can reach  
Senior Technical Editor  
**Brian Dipert**  
at 1-916-760-0159,  
[bdipert@edn.com](mailto:bdipert@edn.com),  
and [www.bdipert.com](http://www.bdipert.com).

